

CHAPTER 11

Pricing Tools in Financial Engineering

1. Introduction

We have thus far proceeded without a discussion of asset pricing *models* and the tools associated with them, as financial engineering has many important dimensions besides pricing. In this chapter, we will discuss models of asset pricing, albeit in a very simple context. A summary chapter on pricing tools would unify some of the previous topics, and show the subtle connections between them. The discussion will approach the issue using a framework that is a natural extension of the financial engineering logic utilized thus far.

Pricing comes with at least *two* problems that seem, at first, difficult to surmount in any satisfactory way. Investors like return, but dislike risk. This means that assets associated with nondiversifiable risks will carry *risk premia*. But, how can we measure such risk premia objectively when buying assets is essentially a matter of *subjective* preferences? Modeling risk premia using *utility functions* may be feasible theoretically, but this is not attractive from a trader's point of view if hundreds of millions of dollars are involved in the process. The potential relationship between risk premia and utility functions of players in the markets is the first unpleasant aspect of practical pricing decisions.

The second problem follows from the first. One way or another, the pricing approach needs to be based on measuring the volatility of future cash flows. But volatility is associated with randomness and with some probability distribution. How can an asset pricing approach that intends to be applicable in practice obtain a reasonable set of real-world probabilities?¹

Modern finance has found an ingenious and *practical* way of dealing with both these questions simultaneously. Instead of using a framework where risk premia are modeled explicitly, the profession *transforms* a problem with risk premia into one where there are no risk premia. Interestingly, this transformation is done in a way that the relevant probability distribution ceases to be the *real-world probability* and, instead, becomes a *market-determined*

¹ Note that the subjective nature of risk premia was in the realm of pure economic theory, whereas the issue of obtaining satisfactory real-world probability distributions falls within the domain of econometric theory.

probability that can be numerically calculated at *any* point in time if there is a reasonable number of liquid instruments.² With this approach, the assets will be priced in an *artificial* risk-neutral environment where the risk premia are indirectly taken into account. This methodology is labeled the *Martingale approach*. It is a powerful tool in practical asset pricing and risk management.

A newcomer to financial engineering may find it hard to believe that a more or less unified theory for pricing financial assets that can be successfully applied in real-world pricing actually exists. After all, there are many different types of assets, and not all of them seem amenable to the same pricing methodology, even at a theoretical level. A market practitioner may already have heard of risk-neutral pricing, but just like the newcomer to financial engineering, he or she may regard the basic theory behind it as very *abstract*. And yet, the theory is surprisingly potent. This chapter provides a discussion of this methodology from the point of view of a financial engineer. Hence, even though the topic is asset pricing, the way we approach it is based on ideas developed in previous chapters. Basically, this pricing methodology is presented as a general approach to synthetic asset creation.

Of course, like any other theory, this methodology depends on some strict assumptions. The methods used in this text will uniformly make one common assumption that needs to be pointed out at the start. Only those models that assume *complete markets* are discussed. In heuristic terms, when markets are “complete,” there are “enough” liquid instruments for obtaining the working probability distribution.

This chapter progressively introduces a number of important theoretical results that are used in pricing, hedging, and risk management application. The *main* result is called the fundamental theorem of asset pricing. Instead of a mathematical proof, we use a *financial engineering argument* to justify it, and a number of important consequences will emerge. Throughout the chapter, we will single out the results that have practical implications.

2. Summary of Pricing Approaches

In this section we remind the reader of some important issues from earlier chapters. Suppose we want to find the *fair* market price of an instrument. First, we construct a synthetic equivalent to this instrument using liquid contracts that trade in financial markets. Clearly, this requires that such contracts are indeed available. Second, once these liquid contracts are found, an arbitrage argument is used. The cost of the *replicating portfolio* should equal the cost of the instrument we are trying to price. Third, a trader would add a proper margin to this cost and thus obtain the fair price.

In earlier chapters, we obtained synthetics for forward rate agreements (FRAs), foreign-exchange (FX) forwards, and several other quasi-linear instruments. Each of these constitutes an early example of asset pricing. Obtain the synthetic and see how much it costs. By adding a profit to this cost, the fair market price is obtained. It turns out that we can *extend* this practical approach and obtain a general theory.

It should be reemphasized that *pricing* and *hedging* efforts can sometimes be regarded as two sides of the same coin. In fact, hedging a product requires finding a replicating portfolio and then using it to cover the position in the original asset. If the trader is long in the original instrument, he or she would be short in the synthetic, and vice versa. This way, exposures to risks would cancel out and the position would become “riskless.” This process results in the creation of a replicating portfolio whose cost cannot be that different from the price of the original asset.

² That is to say, instead of using historical data, we can derive the desired probability distribution from the current quotes.

Thus, a hedge will transfer unwanted risks to other parties but, at the same time, will provide a way to price the original asset.³

Pricing theory is also useful for the creation of “new products.” A new product is basically a series of contingent cash flows. We would, first, put together a combination of financial instruments that have the same cash flows. Then, we would write a separate contract and sell these cash flows to others under a *new name*. For example, a strip of FRAs or futures can be purchased and resulting cash flows are then labeled a swap and sold to others. The new product is, in fact, a dynamically maintained portfolio of existing instruments, and its fair cost will equal the sum of the price of its constituents.

3. The Framework

The pricing framework that we use emphasizes important aspects of the theory within a real-world setting. We assume that m liquid asset prices are observed at times $t_i, i = 1, 2, \dots$. The time t_i price of the k th asset is denoted by S_{kt_i} . The latter can represent credit, stocks, fixed-income instruments, the corresponding derivatives, or commodity prices.

In theory, a typical S_{kt_i} can assume any real value. This makes the number of possible values infinite and *uncountable*. But in practice, every price is quoted to a small number of decimal places and, hence, has a countable number of possible future values. Foreign exchange rates, for example, are in general quoted to four decimal places. This brings us to the next important notion that we would like to introduce.

3.1. States of the World

Let t_0 denote the “present,” and consider the k th asset price S_{kT} , at a future date, $T = t_i$, for some $0 < i$. At time t_0 , the S_{kT} will be a random variable.⁴ Let the symbol ω^j , with $j = 1, \dots, n$ represent time- T *states of the world* that relate to the random variable S_{kT} .⁵ We assume that $n \leq m$, which amounts to saying that there are at least as many liquid assets as there are time- T states of the world. For example, it is common practice in financial markets to assume a “bullish” state, a “bearish” state, and a “no-change” state. Traders expect prices in the future to be either “higher,” “lower,” or to “remain the same.” The ω^j generalizes this characterization, and makes it operational.

EXAMPLE:

In this example, we construct the states of the world that relate to some asset whose time t_i price is denoted by S_{t_i} . Without any loss of generality, let

$$S_{t_0} = 100 \tag{1}$$

Suppose, at a future date T , with $t_n = T$, there are only $n = 4$ states of the world. We consider the task of defining these states.

³ If the hedge is not “perfect,” the market maker will add another margin to the cost to account for any small deviation in the sensitivities toward the underlying risks. For example, if some exotic option cannot be perfectly hedged by the spot and the cash, then the market maker will increase or decrease the price to take into account these imperfections.

⁴ The current value of the asset S_{kt_0} , on the other hand, is known.

⁵ According to this, ω^j may also need a T subscript. But we ignore it and ask the reader to remember this point.

1. Set the value of some grid parameter ΔS to assign neighboring values of S_T into a single state. For example, let

$$\Delta S = 2 \tag{2}$$

2. Next, pick two upper and lower bounds $[S^{\min}, S^{\max}]$ such that the probability of S_T being outside this interval is relatively small and that excursions outside this range can safely be ignored. For example, let $S^{\max} = 104$ and $S^{\min} = 96$. Accordingly, the events $104 < S_T$ and $S_T < 96$ are considered unlikely to occur, and, hence, a detailed breakdown of these states of the world is not needed. Clearly, the choice of numerical values for $[S^{\min}, S^{\max}]$ depends, among other things, on the perceived volatility of S_t during the period $[t_0, T]$.⁶
3. The states of the world can then be defined in the following fashion:

$$\omega^1 = \{S_T \text{ such that } S_T < S^{\min}\} \tag{3}$$

$$\omega^2 = \{S_T \text{ such that } S_T \in [S^{\min}, S^{\min} + \Delta S]\} \tag{4}$$

$$\omega^3 = \{S_T \text{ such that } S_T \in [S^{\min} + \Delta S, S^{\min} + 2\Delta S = S^{\max}]\} \tag{5}$$

$$\omega^4 = \{S_T \text{ such that } S^{\max} < S_T\} \tag{6}$$

This situation is shown in Figure 11-1.

Here, the total number of states of the world depends on the size of the grid parameter ΔS , and on the choice of upper and lower bounds $[S^{\min}, S^{\max}]$. These, in turn, depend on market psychology at time t_0 . For example, selecting the total number of states as $n = 4$ could be justified, if the ranges for S_T shown here were the only ones found relevant for pricing and risk-management problems faced during that particular day. If a problem under consideration requires a *finer* or *coarser* subdivision of the future, the value for n would change accordingly.

3.2. The Payoff Matrix

The next step in obtaining the fundamental theorem of asset pricing is the definition of a payoff matrix for period T . Time- T values of the assets, S_{kt} , depend on the state of the world, ω^i , that will occur at time T . Given that we are working with a finite number of states of the world,




FIGURE 11-1

⁶ In practice, these upper and lower bounds have to be properly *calibrated* to observed liquid, arbitrage-free prices.

possible values for these assets would be easy to list. Let z_k^i represent the value assumed by the k th asset in state ω^i , at time- T :

$$S_{kT}^i = z_k^i \quad (7)$$

Then, for the first n assets, $n \leq m$, we can form the following payoff matrix for time T :

$$D = \begin{pmatrix} z_1^1 & \cdots & z_1^n \\ \vdots & \ddots & \vdots \\ z_n^1 & \cdots & z_n^n \end{pmatrix} \quad (8)$$

A typical row of this matrix would represent possible values of a *particular asset* in different states of the world. A typical column represents different asset prices, in a *particular state of the world*. The definition of ω^i should automatically lead to a definition of the possible values for assets under consideration, as shown in the previous example.

The fundamental theorem of asset pricing is about how “current” asset prices, S_{kt} , relate to the possible values represented by matrix D . We form a matrix equation that will play an important role in the next three chapters.

3.3. The Fundamental Theorem

Consider the *linear* system of equations defined for a series of Q^i , indexed by the state of the world i :

$$\begin{pmatrix} S_{1t_0} \\ \vdots \\ S_{nt_0} \end{pmatrix} = \begin{pmatrix} z_1^1 & \cdots & z_1^n \\ \vdots & \ddots & \vdots \\ z_n^1 & \cdots & z_n^n \end{pmatrix} \begin{pmatrix} Q^1 \\ \vdots \\ Q^n \end{pmatrix} \quad (9)$$

The left-hand side shows the vector of current liquid asset prices observed at time t_0 . The right-hand side has two components. The first is the matrix D of possible values for these prices at time T , and the second is a vector of constants, $\{Q^1, \dots, Q^n\}$. The fundamental theorem of asset pricing concerns this matrix equation and the properties of the $\{Q^i\}$. The theorem can be stated heuristically as follows:

Theorem *The time t_0 prices for the $\{S_{kt_0}\}$ are arbitrage-free if and only if $\{Q^i\}$ exist and are positive.*

Thus, the theorem actually works both ways. If S_{kt_0} are arbitrage-free, then Q^i exist and are all positive. If Q^i exist and are positive, then the S_{kt_0} are arbitrage-free.

The fundamental theorem of asset pricing provides a *unified* pricing tool for pricing real-world assets. In the remaining part of this chapter, we derive important implications of this theorem. These can be regarded as *corollaries* that are exploited routinely in asset pricing. The first of these corollaries is the existence of *synthetic* probabilities. However, before we discuss these results we need to motivate the $\{Q^i\}$ and show why the theorem holds.

3.4. Definition of an Arbitrage Opportunity

What is meant by arbitrage-free prices? To answer this question we need to define *arbitrage opportunity* formally. Formal definition of the framework outlined in this section provides this.

Consider the asset prices S_{1t}, \dots, S_{kt} . Associate the portfolio weights θ_i with asset S_{it} . Then we say that there is an arbitrage opportunity if either of the following two conditions hold.

1. A portfolio with weights θ_i can be found such that:

$$\sum_{i=1}^k \theta_i S_{it} = 0 \quad (10)$$

simultaneously with

$$0 \leq \sum_{i=1}^k \theta_i S_{iT} \quad (11)$$

According to these conditions, the market practitioner advances no cash at time t to form the portfolio, but still has access to some non-zero gains at time T . This is the first type of arbitrage opportunity.

2. A portfolio with weights θ_i can be found such that:

$$\sum_{i=1}^k \theta_i S_{it} \leq 0 \quad (12)$$

simultaneously with

$$\sum_{i=1}^k \theta_i S_{iT} = 0 \quad (13)$$

In this case the market practitioner receives cash at time t while forming the portfolio, but has no liabilities at time T .

It is clear that in either case, the size of these *arbitrage portfolios* is arbitrary since no liabilities are incurred. The formal definition of arbitrage-free prices requires that such portfolios not be feasible at the “current” prices $\{S_{it}\}$.

Notice that what market professionals call an *arbitrage strategy* is very different from this formal definition of arbitrage opportunity. In general, when practitioners talk about “arb” they mean positions that have a relatively small probability of losing money. Clearly this violates both of the conditions mentioned above. The methods introduced in this chapter deal with the lack of formal arbitrage opportunities and not with the market practitioners’ arbitrage strategies. It should be remembered that it is the formal no-arbitrage condition that provides the important tools used in pricing and risk-management.

3.5. Interpreting the Q^i : State Prices

Given the states of the world ω^i , $i = 1, \dots, n$, we can write the preceding matrix equation for the *special case* of two important sets of instruments that are essential to understanding arbitrage-free pricing. Suppose the first asset S_{1t} is a *risk-free savings deposit* account and assume, without any loss of generality, that t_i represents years.⁷ If 1 dollar is deposited at time t_0 , $(1 + r_{t_0})$ can be earned at time t_1 without any risk of *default*. The r_{t_0} is the rate that is observed as of time t_0 .

⁷ This simplifies the notation, since the days’ adjustment parameter will equal one.

The second set of instruments are *elementary insurance contracts*. We denote them by C_i . These contracts are defined in the following way:

- C_1 pays \$1 at time T if ω^1 occurs. Otherwise it pays zero.
- ⋮
- C_n pays \$1 at time T if ω^n occurs. Otherwise it pays zero.

If a market practitioner considers state ω^i as “risky,” he or she can buy a desired number of C_i ’s as insurance to *guarantee* any needed cash flow in that state.

Suppose now that all C_i are actively traded at time t_0 . Then, according to the matrix equation, the correct arbitrage-free prices of these contracts are given by Q^i . This is the case since plugging the current prices of the savings account and the C_i at time t_0 into the matrix equation (9) gives⁸

$$\begin{bmatrix} 1 \\ C_1 \\ \dots \\ C_n \end{bmatrix} = \begin{bmatrix} (1+r_{t_0}) & \dots & \dots & (1+r_{t_0}) \\ 1 & 0 & \dots & 0 \\ \dots & \dots & \dots & \dots \\ 0 & \dots & 0 & 1 \end{bmatrix} \begin{bmatrix} Q^1 \\ \dots \\ Q^n \end{bmatrix} \quad (14)$$

Following from this matrix equation, the Q^i have three important properties. First, we see that they are *equal* to the prices of the elementary insurance contracts:

$$C_i = Q^i \quad (15)$$

It is for this reason that the Q^i are also called *state prices*. Second, we can show that if interest rates are positive, the sum of the time- t_0 prices of C_i is *less than one*. Consider the following: A portfolio that consists of buying *one of each* insurance contract C_i at time t_0 will guarantee 1 dollar at time T no matter which state, ω^i , is realized at time T . But, a guaranteed future dollar should be worth *less* than a current dollar at hand, as long as interest rates are positive. This means that the sum of Q^i paid for the elementary contracts at time t_0 should satisfy

$$Q^1 + Q^2 + \dots + Q^n < 1 \quad (16)$$

From the first row of the matrix equation in (14) we can write

$$\sum_{i=1}^n Q^i (1+r_{t_0}) = 1 \quad (17)$$

After rearranging,

$$Q^1 + Q^2 + \dots + Q^n = \frac{1}{(1+r_{t_0})} \quad (18)$$

The *third* property is a little harder to see. As the fundamental theorem states, none of the Q^i can be *negative* or zero if the C_i are indeed *arbitrage-free*. We show this with a simple counter-example.

EXAMPLE:

Suppose we have $n=4$ and that the first elementary contract has a negative price, $C_1 = -1$. Without any loss of generality, suppose all other elementary insurance contracts have a positive price. Then the portfolio

$$\{(Q^2 + Q^3 + Q^4) \text{ unit of } C_1, 1 \text{ unit of } C_2, 1 \text{ unit of } C_3, 1 \text{ unit of } C_4\} \quad (19)$$

⁸ D in our setup is a $(n+1) \times n$ matrix.

has zero cost at time t_0 , and yet will guarantee a positive return at time t_1 . More precisely, the portfolio returns 1 dollar in states 2 to 4 and $(Q^2 + Q^3 + Q^4)$ dollars in state 1.

Hence, as long as one or more of the Q^i are negative, there will always be an arbitrage opportunity. A trader can “buy” the contract(s) with a negative price and use the cash generated to purchase the other contracts. This way, a positive return at T is *guaranteed*, while at the same time the zero-cost structure of the initial portfolio is maintained. For such arbitrage opportunities not to exist, we need $0 < Q^i$ for all i .

3.5.1. Remarks

Before going further, we ask two questions that may have already troubled the reader given the financial engineering approach we adopted in this book.

- Do insurance contracts such as C_i exist in the real world? Are they actively traded?
- Is the assumption of a small *finite* number of states of the world realistic? How can such a restrictive view of the future be useful in pricing real-world instruments?

In the next few sections we will show that the answer to both of these questions is a qualified yes. To understand this, we need to relate the elementary insurance contracts to the concept of options. Options can be considered as ways of trading baskets of C_i 's. A typical C_i pays 1 dollar if state i occurs and nothing in all other states. Thus, it is clear that option payoffs at expiration are different from those of elementary contracts. Options pay nothing if they expire out-of-the-money, but they pay, $(S_T - K)$ if they expire in-the-money. This means that depending on how we define the states ω^i , unlike the elementary contracts, options can make payments in *more* than one state. But this difference is really not that important since we can get all desired C_i from option prices, if options trade for all strikes K . In other words, the pricing framework that we are discussing here will be much more useful in practice than it seems at first.

As to the second question, it has to be said that, in practice, few strikes of an option series trade actively. This suggests that the finite state assumption may not be that unrealistic after all.

4. An Application

The framework based on state prices and elementary insurance contracts is a surprisingly potent and *realistic* pricing tool. Before going any further and obtaining more results from the fundamental theorem of asset pricing, we prefer to provide a real-world example. The following reading deals with the S&P500 index and its associated options.

EXAMPLE:

The S&P500 is an index of 500 leading stocks from the United States. It is closely monitored by market participants and traded in futures markets. One can buy and sell liquid options written on the S&P500 at the Chicago Board of Options Exchange (CBOE). These options, with an expiration date of December 2001 are shown in Table 11-1 as they were quoted on August 10, 2001.

At the time these data were gathered, the index was at 1187. The three most liquid call options are

$$\{1275 - \text{Call}, 1200 - \text{Call}, 1350 - \text{Call}\} \quad (20)$$

TABLE 11-1

Calls	Last sale	Bid	Ask	Volume	Open interest
Dec 1175	67.1	68	70	51	1378
Dec 1200	46.5	52.8	54.8	150	8570
Dec 1225	41	40.3	42.3	1	6792
Dec 1250	28.5	29.6	31.6	0	11873
Dec 1275	22.8	21.3	23.3	201	6979
Dec 1300	15.8	15	16.2	34	16362
Dec 1325	9.5	10	11	0	9281
Dec 1350	6.8	6.3	7.3	125	8916
Dec 1375	4.1	4	4.7	0	2818
Dec 1400	2.5	2.5	3.2	10	17730
Dec 1425	1.4	1.4	1.85	0	4464
Dec 1450	0.9	0.8	1.25	9	9383
Dec 1475	0.5	0.35	0.8	0	122
Puts	Last sale	Bid	Ask	Volume	Open interest
Dec 800	1.65	1.2	1.65	10	1214
Dec 900	4.3	3.4	4.1	24	11449
Dec 950	5.4	5.3	6.3	10	8349
Dec 995	10.1	8.5	9.5	0	11836
Dec 1025	13	11.1	12.6	11	5614
Dec 1050	13.6	14	15.5	106	19483
Dec 1060	16.5	15.7	17.2	1	1597
Dec 1075	22.5	18	19.5	1	316
Dec 1100	26	22.7	24.7	0	17947
Dec 1150	39	35.3	37.3	2	16587
Dec 1175	44	44.1	46.1	14	4897
Dec 1200	53	53.9	55.9	897	26949

The three most liquid put options, on the other hand, are

$$\{1200\text{-Put}, 1050\text{-Put}, 900\text{-Put}\} \quad (21)$$

Not surprisingly, all the liquid options are out-of-the-money as liquid options generally are.⁹

We will now show how this information can be used to obtain (1) the states of the world ω^i , (2) the state prices Q^i , and (3) the corresponding synthetic probabilities associated with the Q^i . We will do this in the simple setting used thus far.

4.1. Obtaining the ω^i

A financial engineer always operates in response to a particular kind of problem and the states of the world to be defined relative to the needs, *at that time*. In our present example we are working with S&P500 options, which means that the focus is on equity markets. Hence, the corresponding states of the world would relate to different states in which the U.S. stock market

⁹ Buying and selling in-the-money options does not make much sense for market professionals. Practitioners carry these options by borrowing the necessary funds and hedge them immediately. Hence, any intrinsic value is offset by the hedge side anyway. Yet, the convexity of in-the-money options will be the same as with those that are out-of-the-money.

might be at a future date. Also, we need to take into account that trader behavior singles out a relatively small number of *liquid* options with expirations of about three months. For the following example refer back to Table 11-1.

EXAMPLE:

We let S_T represent the value of the S&P500 at expiration and then use the strike prices K_i of the liquid options to define the future states of the world. In fact, strike prices of puts and calls discussed in the preceding example divide the S_T -axis into intervals of equal length. But only a handful of these options are liquid, implying that fine subdivisions were perhaps not needed by the markets for that day and that particular expiration. Accordingly, we can use the strike prices of the three liquid out-of-the-money puts to obtain the intervals

$$\omega^1 = S_T < 900 \quad (22)$$

$$\omega^2 = 900 \leq S_T < 1050 \quad (23)$$

$$\omega^3 = 1050 \leq S_T < 1200 \quad (24)$$

Note that the liquid puts lead to intervals of equal length. It is interesting, but also expected, that the liquid options have this kind of regularity in their strikes. Next, we use the three out-of-the-money calls to get three intervals to define three additional states of the world as

$$\omega^4 = 1200 \leq S_T < 1275 \quad (25)$$

$$\omega^5 = 1275 \leq S_T < 1350 \quad (26)$$

$$\omega^6 = 1350 \leq S_T \quad (27)$$

Here, the last interval is obtained from the highest-strike liquid call option. Figure 11-2 shows these options and the implied intervals. Since these intervals relate to future values of S_T , we consider them the relevant states of the world for S_T .

We pick the midpoint of the bounded intervals as an approximation to that particular state. Let these midpoints be denoted by $\{\bar{S}^i, i = 2, \dots, 5\}$. These midpoints can then be used as a finite set of points that represent ω^i . For the first and last half-open intervals, we select the values of the two extreme points, \bar{S}^1 , and \bar{S}^6 , arbitrarily for the time being. We let

$$\bar{S}^1 = 750 \quad (28)$$

$$\bar{S}^6 = 1400 \quad (29)$$

so that the distance between \bar{S}^i remains constant. This arbitrary selection of the “end states” is clearly unsatisfactory. In fact, by doing this we are in a sense setting the volatility of the random variables arbitrarily. We can, however calibrate our selection. Once our educated guesses are plugged in, we can try to adjust these extreme values so that the resulting Q^i all become positive and price some other liquid asset correctly. In a sense, calibration is an attempt to see which value of the two “end states” replicates the observed prices. But for the time being, we ignore this issue and assume that the end points are selected correctly.

It is open to debate if selecting just six states of the world, as in the example, might represent future possibilities concerning S_T accurately. Traders dealing with the risk in the example must

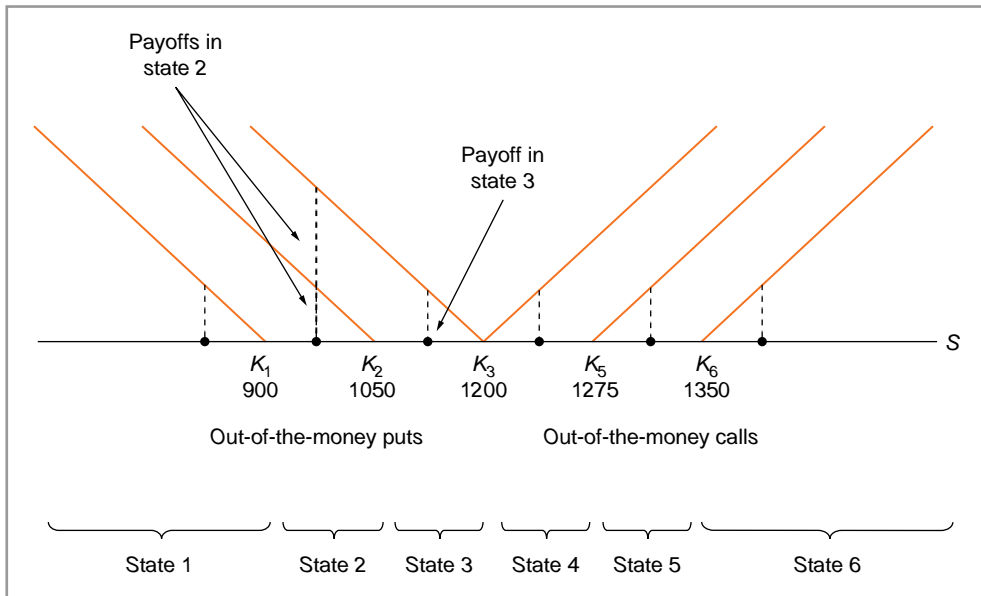


FIGURE 11-2

have thought so, since on that particular date trading approximately *six* liquid options was sufficient to resolve their tasks. It seems that if a *finer* subdivision of the future possibilities were more appropriate, then more liquid strikes would have been traded.

Hence, as usual in financial engineering, the specific values of ω^i that we select are based on the values of liquid instruments. In our case, the possible states of the world were chosen as dictated by liquid call and put options.

4.2. Elementary Contracts and Options

Elementary insurance contracts C_i do not trade directly in world financial markets. Yet, the C_i are not far from a well-known instrument class—options—and they trade “indirectly.” This section shows how elementary insurance contracts can be obtained from options, and vice versa. Plain vanilla options are, in fact, close relatives of elementary insurance contracts. The best way to see this is to consider a numerical example. (Generalizations are straightforward.)

EXAMPLE:

Start with the first and the last options selected for the previous example. Note that the 900-put is equivalent to $K_1 - \bar{S}^1$ units of C_1 because it pays approximately this many dollars if state ω_1 occurs and nothing in all other states. Similarly, the 1350-call is equivalent to $\bar{S}^6 - K_6$ units of C_6 because it pays approximately this many dollars if state 6 occurs and nothing otherwise.

The other calls and puts have payoffs in more than one state, but they also relate to elementary contracts in straightforward ways. For example, the 1050-put is equivalent to a portfolio of two elementary insurance contracts, $K_2 - \bar{S}^1$ units of C_1 and $K_2 - \bar{S}^2$ units of C_2 , because it makes these payments in states 1 and 2, respectively, and nothing else in other states. In fact, pursuing this reasoning, we can obtain the following matrix

equation between the payoffs of elementary contracts C_1, \dots, C_6 and the option prices:

$$\begin{bmatrix} 900\text{-Put} \\ 1050\text{-Put} \\ 1200\text{-Put} \\ 1200\text{-Call} \\ 1275\text{-Call} \\ 1350\text{-Call} \end{bmatrix} = \begin{bmatrix} z_1^1 & 0 & 0 & 0 & 0 & 0 \\ z_2^1 & z_2^2 & 0 & 0 & 0 & 0 \\ z_3^1 & z_3^2 & z_3^3 & 0 & 0 & 0 \\ 0 & 0 & 0 & z_4^4 & z_4^5 & z_4^6 \\ 0 & 0 & 0 & 0 & z_5^5 & z_5^6 \\ 0 & 0 & 0 & 0 & 0 & z_6^6 \end{bmatrix} \begin{bmatrix} C_1 \\ C_2 \\ C_3 \\ C_4 \\ C_5 \\ C_6 \end{bmatrix} \quad (30)$$

This equation holds since we have

$$Q^i = C_i \quad (31)$$

for all i .¹⁰

Thus, given the arbitrage-free values of traded puts and calls with different strikes but similar in every other aspect, we can easily obtain the values of the elementary insurance contracts C_i by inverting the (6×6) matrix on the right side. In fact, it is interesting to see that the matrix equation in the example contains two triangular subsystems that can be solved separately and recursively.

Hence, the existence of liquid options makes a direct application of the fundamental theorem of asset pricing possible. Given a large enough number of liquid option contracts, we can obtain the state prices, Q^i , if these exist, and, if they are all positive, use them to price other illiquid assets that depend on the same risk.¹¹ Obviously, when traders deal with interest rate, or exchange rate risk, or when they are interested in pricing contracts on commodities, they would use liquid options for *those* particular sectors and work with different definitions of the state of the world.

4.3. Elementary Contracts and Replication

We now show how elementary insurance contracts and options that belong to a series can be used in replicating instruments with *arbitrary* payoffs. Consider an arbitrary financial asset, S_t , that is worth z_T^i in state of the world ω^i , $i = 1, \dots, n$, at time T . Given n elementary insurance contracts C_i , we can immediately form a *replicating portfolio* for this asset. Assuming, without any loss of generality, that the time- T payoffs of the S_t asset are denoted by $0 < z_T^i$, we can consider buying the following portfolio:

$$\{z_T^1 \text{ units of } C_1, z_T^2 \text{ units of } C_2, \dots, z_T^n \text{ units of } C_n\} \quad (32)$$

At time T , this portfolio should be worth exactly the same as the S_t , since whatever state occurs, the basket of insurance contracts will make the same time- T payoff as the original asset. This provides an immediate *synthetic* for the S_t . Accordingly, if there are no arbitrage opportunities, the value of the portfolio and the value of the S_t will be identical as of time t as well.¹² We consider an example.

¹⁰ If one needs to get the values of C_i from the traded puts and calls, one should start with the first put, then move to the second put, then the third. The same strategy can be repeated with the last call and so on.

¹¹ Here, this risk is the S&P500 index.

¹² As usual, it is assumed that the S_t makes no other interim payouts.

EXAMPLE:

Take any four independent assets S_{kt} , $k = 1, \dots, 4$ with different payoffs, z_k^i , in the states $\{\omega^i, i = 1, \dots, 4\}$. We can express each one of these assets in terms of the elementary insurance contracts. In other words, we can find one synthetic for each S_{kt} by purchasing the portfolios:

$$\{z_k^1 \text{ unit of } C_1, z_k^2 \text{ unit of } C_2, z_k^3 \text{ unit of } C_3, z_k^4 \text{ unit of } C_4\} \quad (33)$$

Putting these in matrix form, we see that arbitrage-free values, S_{kt_0} , of these assets at time t_0 have to satisfy the matrix equation:

$$\begin{bmatrix} 1 \\ S_{1t_0} \\ S_{2t_0} \\ S_{3t_0} \\ S_{4t_0} \end{bmatrix} = \begin{bmatrix} 1+r_{t_0} & 1+r_{t_0} & 1+r_{t_0} & 1+r_{t_0} \\ z_1^1 & z_1^2 & z_1^3 & z_1^4 \\ z_2^1 & z_2^2 & z_2^3 & z_2^4 \\ z_3^1 & z_3^2 & z_3^3 & z_3^4 \\ z_4^1 & z_4^2 & z_4^3 & z_4^4 \end{bmatrix} \begin{bmatrix} Q^1 \\ Q^2 \\ Q^3 \\ Q^4 \end{bmatrix} \quad (34)$$

where the matrix on the right-hand side contains all possible values of the assets S_{kt} in states ω^i , at time T .¹³

Hence, given the prices of actively traded elementary contracts C_i , we can easily calculate the time- t cost of forming the portfolio:

$$\text{Cost} = C_1 z_T^1 + C_2 z_T^2 + \dots + C_n z_T^n \quad (35)$$

This can be regarded as the cost basis for the S_t asset. Adding a proper margin to it will give the fair market price S_t .

EXAMPLE:

Suppose the S_t -asset has the following payoffs in the states of the world $i = 1, \dots, 4$:

$$\{z_T^1 = 10, z_T^2 = 1, z_T^3 = 14, z_T^4 = 16\} \quad (36)$$

Then, buying 10 units of the first insurance contract C_1 will guarantee the 10 in the first state, and so on.

Suppose we observe the following prices for the elementary insurance contracts:

$$C_1 = .3, C_2 = .2, C_3 = .4, C_4 = .07 \quad (37)$$

Then the total cost of the insurance contracts purchased will be

$$\text{Cost} = (.3)10 + (.2)(1) + (.4)14 + (.07)16 \quad (38)$$

$$= 9.92 \quad (39)$$

This should equal the current price of S_t once a proper profit margin is added.

¹³ For notational simplicity, we eliminated time subscripts in the matrix equation. Also remember that the time index represents, years and that $C_i = Q^i$.

Clearly, if such elementary insurance contracts actively traded in financial markets, the job of a financial engineer would be greatly simplified. It would be straightforward to construct synthetics for *any* asset, and then price them using the cost of the replicating portfolios as shown in the example. However, there is a close connection between C_i and options of the same series that differ only in their strikes. We saw how to obtain the C_i from liquid option prices. Accordingly, if options with a broad array of strikes trade in financial markets, then traders can create *static* replicating portfolios for assets with arbitrary payoffs.¹⁴

5. Implications of the Fundamental Theorem

The fundamental theorem of asset pricing has a number of implications that play a critical role in financial engineering and derivatives pricing. First, using this theorem we can obtain probability distributions that can be used in asset pricing. These probability distributions will be objective and operational. Second, the theorem leads to the so-called Martingale representation of asset prices. Such a representation is useful in modeling asset price dynamics. Third, we will see that the Martingale representation can serve to objectively set expected asset returns. This property eliminates the need to model and estimate the “drift factors” in asset price dynamics. We will now study these issues in more detail.

5.1. Result 1: Risk-Adjusted Probabilities

The Q^i introduced in the previous section can be modified judiciously in order to obtain convenient probability distributions that the financial engineer can work with. These distributions do *not* provide real-world odds on the states of the world ω^i , and hence cannot be used directly in econometric prediction. Yet they *do* yield correct arbitrage-free prices. (This section shows how.) But there is more. As there are *many* such distributions, the market practitioner can also *choose* the distribution that fits his/her current needs best. How to make this choice is discussed in the next section.

5.1.1. Risk-Neutral Probabilities

Using the state prices Q^i , we first obtain the so-called *risk-neutral probability* distribution. Consider the first row of the matrix equation (9). Assume that it represents the savings account.¹⁵

$$(1 + r_{t_0})Q^1 + \cdots + (1 + r_{t_0})Q^n = 1 \quad (40)$$

Relabel, using

$$\tilde{p}_i = (1 + r_{t_0})Q^i \quad (41)$$

According to equation (40), we have

$$\tilde{p}_i + \cdots + \tilde{p}_n = 1 \quad (42)$$

where

$$0 < \tilde{p}_i \quad (43)$$

¹⁴ Again, the asset with arbitrary payoffs should depend on the same underlying risk as the options.

¹⁵ Remember that $T - t_0$ is assumed to be one year.

since each Q^i is positive. This implies that the numbers \tilde{p}_i have the properties of a discrete *probability distribution*. They are positive and they add up to one. Since they are determined by the markets, we call them “risk-adjusted” probabilities. They are, in fact, obtained as linear combinations of n current asset prices. This particular set of synthetic probabilities is referred to as the *risk-neutral* probability distribution.

To be more precise, the risk-neutral probabilities $\{\tilde{p}_i\}$ are time- t_0 probabilities on the states that occur at time T . Thus, if we wanted to be more exact, they would have to carry two more subscripts, t_0 and T . Yet, these will be omitted for notational convenience and assumed to be understood by the reader.

5.1.2. Other Probabilities

Several other synthetic probabilities can be generated, and these may turn out to be more useful than the risk-neutral probabilities. Given the positive Q^i , we can rescale these by any positive normalizing factor so that they can be interpreted as probabilities. There are many ways to proceed. In fact, *as long as a current asset price S_{kt_0} is nonzero and the z_k^j are positive*, one can choose any k th row of the matrix equation in (9) to write

$$1 = \sum_{i=1}^n \frac{z_k^i}{S_{kt_0}} Q^i \quad (44)$$

and then define

$$\frac{z_k^i}{S_{kt_0}} Q^i = \tilde{p}_i^k \quad (45)$$

The $\tilde{p}_i^k, i = 1, \dots, n$ can be interpreted as probabilities obtained after *normalizing* by the S_{kt_0} asset. The \tilde{p}_i^k will be positive and will add up to one. Hence, they will have the characteristics of a probability distribution, but again they cannot be used in prediction since they are not the actual probabilities of a particular state of the world ω^i occurring. Clearly, for *each* nonzero S_{kt_0} we can obtain a new probability \tilde{p}_i^k . These will be different across states ω^i , as long as the time- T value of the asset is positive in all states.¹⁶

It turns out that *how* we normalize a sequence of $\{Q^i\}$ in order to convert them into some synthetic probability is important. The special case, where

$$S_{kt_0} = B(t_0, T) \quad (46)$$

$B(t_0, T)$ being the current price of a T -maturity risk-free pure discount bond, is especially interesting. This yields the so-called *T-forward measure*. Because the discount bond matures at time T , the time- T values of the asset are given by

$$z_k^i = 1 \quad (47)$$

for all i .

Thus, we can simply divide state prices Q^i by the current price of a default-free discount bond maturing at time t_0 , and obtain the T -forward measure:

$$\tilde{p}_i^T = Q^i \frac{1}{B(t_0, T)} \quad (48)$$

We will see in Chapter 13 that the T -forward measure is the natural way to deal with payoffs associated with time T . Let's consider an example.

¹⁶ We emphasize that we need a positive price and positive possible time- T values in order to do this.

EXAMPLE:

Suppose short-term risk-free rates are 5% and that there are four states of the world. We observe the following arbitrage-free bid prices for four assets at time t_0 :

$$S_{1t_0} = 2.45238, S_{2t_0} = 1.72238, S_{3t_0} = 6.69429, S_{4t_0} = 3.065 \quad (49)$$

It is assumed that at time $T = t_0 + 1$, measured in years, the four possible values for each asset will be given by the matrix:

$$\begin{bmatrix} 10 & 3 & 1 & 1 \\ 2 & 3 & 2 & 1 \\ 1 & 10 & 10 & 1 \\ 8 & 2 & 10 & 2 \end{bmatrix} \quad (50)$$

We can form the matrix equation and then solve for the corresponding Q^i :

$$\begin{bmatrix} 1 \\ S_{1t_0} \\ S_{2t_0} \\ S_{3t_0} \\ S_{4t_0} \end{bmatrix} = \begin{bmatrix} 1 + .05 & 1 + .05 & 1 + .05 & 1 + .05 \\ 10 & 3 & 1 & 1 \\ 2 & 3 & 2 & 1 \\ 1 & 10 & 10 & 6 \\ 8 & 2 & 10 & 2 \end{bmatrix} \begin{bmatrix} Q^1 \\ Q^2 \\ Q^3 \\ Q^4 \end{bmatrix} \quad (51)$$

Using the first four rows of this system, we solve for the Q^i :

$$Q^1 = 0.1, Q^2 = 0.3, Q^3 = 0.07, Q^4 = 0.482 \quad (52)$$

Next, we obtain the risk-neutral probabilities by using

$$\tilde{p}_i = (1 + .05)Q^i \quad (53)$$

which gives

$$\tilde{p}_1 = 0.105, \tilde{p}_2 = 0.315, \tilde{p}_3 = 0.0735, \tilde{p}_4 = 0.5065 \quad (54)$$

As a final point, note that we used the first four rows of the system shown here to determine the values of Q^i . However, the price of S_{4t_0} is also arbitrage-free:

$$\sum_{i=1}^4 Q^i S_{4T}^i = 3.065 \quad (55)$$

as required.

Interestingly, for short-term instruments, and with “normal” short-term interest rates of around 3 to 5%, the savings account normalization makes little difference. If $T - t_0$ is small, the Q^i will be only marginally different from the \tilde{p}_i .¹⁷

¹⁷ For example, at a 5% short rate, a 1-month discount bond will sell for approximately

$$B(t_0, t) = \frac{1}{1 + .05 \frac{1}{12}} = .9958$$

so that dividing by this scale factor will not modify the Q^i very much.

5.1.3. A Remark

Can derivatives be used for normalization? For example, instead of normalizing by a savings account or by using bonds, could we normalize with a swap? The answer is no. There are probabilities called swap measures, but the normalization that applies in these cases is not a swap, but an annuity. Most derivatives are not usable in the normalization process because normalization by an S_{kt} implies, essentially, that the state prices Q^i are multiplied by factors such as

$$\frac{z_k^i S_{kt}}{S_{kt} z_k^i} = 1 \quad (56)$$

and then grouped according to

$$\frac{z_k^i S_{kt}}{S_{kt} z_k^i} Q^i = \frac{S_{kt}}{z_k^i} \tilde{p}_i^k \quad (57)$$

But in this operation, both the z_k^i , $i = 1, \dots, n$ and the S_{kt} should be nonzero. Otherwise the ratios would be undefined. This will be seen below.

5.1.4. Swap Measure

The normalizations thus far used only *one* asset, S_{kt} , in converting the Q^i into probabilities \tilde{p}^k . This need not be so. We can normalize using a linear combination of many assets, and sometimes this proves very useful. This is the case for the so-called swap, or annuity, measure. The swap measure is dealt with in Chapter 21.

5.2. Result 2: Martingale Property

The fundamental theorem of asset pricing also provides a convenient model for pricing and risk-management purposes. All *properly normalized* asset prices have a *Martingale property* under a properly selected synthetic probability \tilde{P}^k . Let X_t be a stochastic process that has the following property:

$$X_t = E_t^{\tilde{P}^k} [X_T] \quad t < T \quad (58)$$

This essentially says that the X_t have no predictable trend for all t . X_t is referred to as a Martingale. To see how this can be applied to asset pricing theory, first choose the risk-neutral probability \tilde{P} as the working probability distribution.

5.2.1. Martingales under \tilde{P}

Consider any k th row in the matrix equation (9)

$$S_{kt_0} = (z_k^1)Q^1 + (z_k^2)Q^2 + \dots + (z_k^n)Q^n \quad (59)$$

Replace the Q^i with the risk-neutral probabilities \tilde{p}_i using

$$Q^i = \frac{1}{1 + r_{t_0}} \tilde{p}_i \quad (60)$$

r_{t_0} is the interest on a risk-free one-year deposit.

This gives

$$S_{kt_0} = \frac{1}{1 + r_{t_0}} [z_k^1 \tilde{p}_1 + \cdots + z_k^n \tilde{p}_n] \quad (61)$$

Here, the right-hand side is an *average* of the future values of S_{kT} , weighted by \tilde{p}_i . Thus, bringing back the time subscripts, the current arbitrage-free price S_{kt_0} satisfies

$$S_{kt_0} = \frac{1}{(1 + r_{t_0})} E_{t_0}^{\tilde{P}}[S_{kT}] \quad t_0 < T \quad (62)$$

In general terms, letting

$$X_t = \frac{\text{Time-}t \text{ value of } S_{kt}}{\text{Time-}t \text{ value of the savings account}} \quad (63)$$

we see that asset values normalized by the savings deposit have the Martingale property:

$$X_t = E_t^{\tilde{P}}[X_T] \quad t < T \quad (64)$$

Thus, all tools associated with Martingales immediately become available to the financial engineer for pricing and risk management.

5.2.2. Martingales under Other Probabilities

The convenience of working with Martingales is not limited to the risk-neutral measure \tilde{P} . A normalization with any nonzero price S_{jt} will lead to another Martingale. Consider the same k th row of the matrix equation in (9)

$$S_{kt_0} = (z_k^1)Q^1 + \cdots + (z_k^n)Q^n \quad (65)$$

This time, replace the Q^i using the $S_{jt_0}, j \neq k$, normalization:

$$\tilde{p}_i^j = Q^i \frac{z_j^i}{S_{jt_0}} \quad (66)$$

We obtain, assuming that the denominator elements are positive:

$$S_{kt_0} = S_{jt_0} \left[z_k^1 \frac{1}{z_j^1} \tilde{p}_1^j + \cdots + z_k^n \frac{1}{z_j^n} \tilde{p}_n^j \right] \quad (67)$$

this means that the ratio,

$$X_t = \frac{S_{kt}}{S_{jt}} \quad (68)$$

is a Martingale under the \tilde{P}^j measure:

$$X_t = E_t^{\tilde{P}^j}[X_T] \quad t < T \quad (69)$$

It is obvious that the probability associated with a particular Martingale is a function of the normalization that is chosen, and that the implied Martingale property can be exploited in pricing. By choosing a Martingale, the financial engineer is also choosing the probability that he or she will be *working* with. In the remainder of this chapter and in the next, we will see several examples of how Martingale properties can be utilized.

5.3. Result 3: Expected Returns

The next implication of the fundamental theorem is useful in modeling arbitrage-free *dynamics* for asset prices. Every synthetic probability leads to a particular *expected return* for the asset prices under consideration. These expected returns will *differ* from the true (subjective) expectations of players in the markets, but because they are agreed upon by all market participants and are associated with arbitrage-free prices, they will be even more useful than the true expectations.

We conduct the discussion in terms of the risk-neutral probability \tilde{P} , but our conclusions are valid for all other \tilde{P}^k . Consider again the Martingale property for an asset whose price is denoted by S_t , but this time reintroduce the day's adjustment parameter δ , dropping the assumption that t_i represents years. We can write, for some $0 < \delta$,

$$S_t = \frac{1}{(1 + r_t\delta)} E_t^{\tilde{P}}[S_{t+\delta}] \quad (70)$$

Rearrange to obtain

$$(1 + r_t\delta) = E_t^{\tilde{P}} \left[\frac{S_{t+\delta}}{S_t} \right] \quad (71)$$

According to this expression, under the probability \tilde{P} , expected net annual returns for *all* liquid assets will equal r_t , the risk-free rate observed at time t .

Similar results concerning the expected returns of the assets are obtained under other probabilities \tilde{P}^k . The expected returns will be different under different probabilities. Market practitioners can select the working probability so as to set the expected return of the asset to a *desired* number.¹⁸

In Chapter 13, we will see more complicated applications of this idea using time- T forward measures. There, the expected change in the forward rates is set equal to zero by a judicious choice of probabilities.

5.3.1. Martingales and Risk Premia

Let us see how the use of Martingales “internalizes” the risk premia associated with nondiversifiable market risks. Let $X_t, t \in [t_0, T]$ be a risky asset and $\Delta > 0$ be a small time interval. The annualized *gross return* of the X_t as expected by *players* at time t , is defined by

$$1 + \hat{R}_t\Delta = E_t^P \left[\frac{X_{t+\Delta}}{X_t} \right] \quad (72)$$

where P represents the *real-world probability* used by market participants in setting up their expectation. Since this is an actual market expectation, the gross return will contain a risk premium:

$$\hat{R}_t = r_t + \mu_t \quad (73)$$

where r_t is the risk-free rate, and μ_t is the *risk premium* commanded by the risky asset.¹⁹ Putting these together, we have

$$(1 + r_t\Delta + \mu_t\Delta) = E_t^P \left[\frac{X_{t+\Delta}}{X_t} \right] \quad (74)$$

¹⁸ Consequently, the associated risk premium need not be estimated.

¹⁹ Under *rational* expectations, the subjective probability P is the same as the “true” distribution of X_t .

or

$$X_t = \frac{1}{(1 + r_t\Delta + \mu_t\Delta)} E_t^P[X_{t+\Delta}] \quad (75)$$

This equality states that the asset price $X_{t+\Delta}$ *discounted* by the factor $(1 + r_t\Delta + \mu_t\Delta)$ is a Martingale *only* if we use the probability P . Note that in this setup there are *two* unknowns: (1) the risk premium μ_t , and (2) the real-world probability P .²⁰ Future cash flows accordingly need to be discounted by subjective discount factors and real-world probabilities need to be estimated. The pricing problem under these conditions is more complex. Financial engineers have to determine the value of the risk premium in addition to “projecting” future earnings or cash flows.

Now consider an alternative. Setting the (positive) risk premium equal to zero in the previous equation gives the inequality

$$X_t < \frac{1}{(1 + r_t\Delta)} E_t^P[X_{t+\Delta}] \quad (76)$$

But this is the same as risk-free savings account normalization. This means that by switching from P to \tilde{P} , we can restore the equality

$$X_t = \frac{1}{(1 + r_t\Delta)} E_t^{\tilde{P}}[X_{t+\Delta}] \quad (77)$$

Thus, normalization and synthetic probabilities internalize the risk premia by converting both unknowns into a known and objective probability \tilde{P} . Equation (77) can be exploited for pricing and risk management.

6. Arbitrage-Free Dynamics

The last result that we derive from the fundamental theorem of asset pricing is a combination of all the corollaries discussed thus far. The synthetic probabilities and the Martingale property that we obtained earlier can be used to derive several *arbitrage-free dynamics* for an asset price. These arbitrage-free dynamics play an important role in pricing situations where an exact synthetic cannot be created, either due to differences in nonlinearities, or due to a lack of liquid constituent assets. In fact, most of the pricing models will proceed along the lines of first obtaining arbitrage-free dynamics, and then either simulating paths from this or obtaining the implied binomial or trinomial trees. PDE methods also use these arbitrage-free dynamics.

6.1. Arbitrage-Free SDEs

In this section we briefly discuss the use of stochastic differential equations as a tool in financial engineering and then show how the fundamental theorem helps in specifying explicit SDEs that can be used in pricing and hedging in practice.²¹ Consider an asset price S_t . Suppose we divide the time period $[t, T]$ into small intervals of equal size Δ . For each time $t + i\Delta$, $i = 1, \dots, n$, we observe a different $S_{t+i\Delta}$. The $S_{t+\Delta} - S_t$ is the *change* in asset price at time t . Choose a working probability from all available synthetic probabilities, and denote it by P^* .

²⁰ Although this latter is *estimable* using econometric methods.

²¹ Appendix 8-2 in Chapter 8 provided the definition and some motivation for SDEs.

Then, we can always calculate the expected value of this change under this probability. In the case of $P^* = \tilde{P}$, we obtain the risk-neutral expected net return by

$$E_t^{\tilde{P}}[S_{t+\Delta} - S_t] = r_t S_t \Delta \quad (78)$$

Next, note that the following statement is *always* true:

$$\text{Actual change in } S_t = \text{“Expected” change} + \text{“Unexpected” change} \quad (79)$$

Now we can use the probability switching method and exploit the Martingale property. For example, for the risk-neutral probability we have

$$[S_{t+\Delta} - S_t] = E_t^{\tilde{P}}[S_{t+\Delta} - S_t] + \epsilon_t \quad (80)$$

where the ϵ_t represents a random variable with zero expectation under the \tilde{P} . Replace from (78)

$$[S_{t+\Delta} - S_t] = r_t S_t \Delta + \epsilon_t \quad (81)$$

Now the error term ϵ_t can be written in the equivalent form

$$\epsilon_t = \sigma(S_t) S_t \Delta W_t \quad (82)$$

where the ΔW_t is a Wiener process increment with variance equal to Δ .

Thus, the arbitrage-free dynamics under the \tilde{P} measure can be written as

$$[S_{t+\Delta} - S_t] = r_t S_t \Delta + \sigma(S_t) S_t \Delta W_t \quad (83)$$

Letting $\Delta \rightarrow 0$, this equation becomes a stochastic differential equation (SDE), that represents the arbitrage-free dynamics under the synthetic probability, \tilde{P} , during an infinitesimally short period dt . Symbolically, the SDE is written as

$$dS_t = r_t S_t dt + \sigma(S_t) S_t dW_t \quad (84)$$

The dS_t and dW_t represent changes in the relevant variables during an *infinitesimal* time interval. Given the values for the (percentage) volatility parameter, $\sigma(S_t)$, these equations can be used to generate arbitrage-free trajectories for the S_t . We deal with these in the next chapter. Note a major advantage of using the risk-neutral probability. The drift term, that is to say the first term on the right-hand side, is known. At this point we consider a second way of obtaining arbitrage-free paths.

6.2. Tree Models

We will see another major application of the Martingale property. We develop the notion of binomial (trinomial) trees introduced in Chapter 7 and obtain an alternative way of handling arbitrage-free dynamics. Suppose the dynamics of S_t can be described by a (geometric) SDE:

$$dS_t = r S_t dt + \sigma S_t dW_t \quad (85)$$

where the volatility is assumed to be given by

$$\sigma(S_t) = \sigma \quad (86)$$

This is the constant *percentage* volatility of S_t . Also note that r_t is set to a constant. It can be shown that this stochastic differential equation can be “solved” for S_t to obtain the relationship (for example, see Øksendal (2003)).

$$S_{t+\Delta} = S_t e^{r\Delta - \frac{1}{2}\sigma^2\Delta + \sigma(W_{t+\Delta} - W_t)} \tag{87}$$

Our purpose is to construct an approximation to the arbitrage-free dynamics of this S_t . We will do this by considering *approximations* to possible paths that S_t can follow between t and some “expiration” date T . This approximation will be such that S_t will satisfy the Martingale property under a judiciously chosen probability. Finally, the approximation should be chosen so that as $\Delta \rightarrow 0$, the mean and the variance of the discrete approximation converge to those of the continuous time process under the relevant probability. It turns out that this can be done in *many* different ways. Each method may have its advantages and disadvantages. We discuss two different ways of building trees. As $\Delta \rightarrow 0$, the dynamics become those of continuous time.

6.2.1. Case 1

The method introduced by Cox-Ross-Rubinstein (CRR) selects the following approximation. First, the period $[t_0, T]$ is divided into N subintervals of equal length. Then, it is assumed that at each point of a path there are possible states. In the CRR case, $n = 2$ and the paths become *binomial*. An alternative trinomial tree is shown in Figure 11-3.




FIGURE 11-3

- At every node i of a possible path, there are only two possible states represented by the numbers $\{u_i, d_i\}$, with the (marginal) probabilities p and $(1 - p)$. The dynamics are selected as follows:

$$S_i^u = u_i S_{i-\Delta} \quad (88)$$

$$S_i^d = d_i S_{i-\Delta} \quad (89)$$

where S_i is the shortcut notation for $S_{t+i\Delta}$.

- The $\{u_i, d_i\}$ are assumed to be constant at u, d .

We now show how to determine the Martingale probabilities. One approach is to find probabilities such that under $p, (1 - p)$:

$$S_i = e^{-r\Delta} E_i^{\tilde{P}}[S_{i+\Delta}] \quad (90)$$

or

$$S_i = e^{-r\Delta} [p S_{i+\Delta}^u + (1 - p) S_{i+\Delta}^d] \quad (91)$$

Using the definition of $S_{i+\Delta}^u, S_{i+\Delta}^d$, in equations (88) and (89), we can write

$$S_i = e^{-r\Delta} [p S_i u + (1 - p) S_i d] \quad (92)$$

The mean and the variance of S_i under this probability should also be as given by the postulated dynamics of the continuous time process in the limit.²² In other words, the p should also satisfy

$$E_i^{\tilde{P}}[S_{i+\Delta}] = [pu + (1 - p)d] S_i \quad (93)$$

and

$$E_i^{\tilde{P}}[S_{i+\Delta}^2 - E_i^{\tilde{P}}[S_{i+\Delta}]^2] = [pu^2 + (1 - p)d^2] S_i^2 - E_i^{\tilde{P}}[S_{i+\Delta}]^2 \quad (94)$$

Use

$$E_i^{\tilde{P}}[S_{i+\Delta}] = S_i e^{r\Delta} \quad (95)$$

$$E_i^{\tilde{P}}[S_{i+\Delta}^2 - E_i^{\tilde{P}}[S_{i+\Delta}]^2] = S_i^2 e^{2r\Delta} (e^{\sigma^2 \Delta} - 1) \quad (96)$$

and get the equations

$$e^{r\Delta} = pu + (1 - p)d \quad (97)$$

$$e^{2r\Delta + \sigma^2 \Delta} = pu^2 + (1 - p)d^2 \quad (98)$$

The p, u, d that satisfy these two equations will (1) satisfy the Martingale equality for all Δ , (2) get arbitrarily close to the mean and the variance of the continuous time process S_t as Δ goes to zero, and (3) make the asymptotic distribution of S_i normal. However, there is one problem. Note that here we have two equations and three unknowns: u, d , and p . One more equation is needed. Choose

$$u = \frac{1}{d} \quad (99)$$

²² Here the probability \tilde{P} is represented by the parameter P .

This makes the tree recombine and completes the system of equations. Under these conditions, the following values solve the equations

$$p = \frac{e^{r\Delta} - d}{u - d} \tag{100}$$

$$u = e^{\sigma\sqrt{\Delta}} \tag{101}$$

$$d = e^{-\sigma\sqrt{\Delta}} \tag{102}$$

Any approximation here is in the sense that all terms containing higher orders of Δ are ignored.²³

6.2.2. Case 2

The previous selection of p, u, d satisfies

$$S_i = e^{-r\Delta} [pS_i e^{\sigma\sqrt{\Delta}} + (1 - p)S_i e^{-\sigma\sqrt{\Delta}}] \tag{103}$$

It turns out that p, u, d can be selected in other ways as well. In particular, note that during an interval Δ , the S_t moves to

$$S_{t+\Delta} = S_t e^{r\Delta - \frac{1}{2}\sigma^2\Delta + \sigma[W_{t+\Delta} - W_t]} \tag{104}$$

Using the approximation

$$W_{t+\Delta} - W_t = \begin{cases} +\sqrt{\Delta} & \text{with probability } .5 \\ -\sqrt{\Delta} & \text{with probability } .5 \end{cases} \tag{105}$$

we get new values for p, u , and d :

$$u = e^{r\Delta - \frac{1}{2}\sigma^2\Delta + \sigma\sqrt{\Delta}} \tag{106}$$

$$d = e^{r\Delta - \frac{1}{2}\sigma^2\Delta - \sigma\sqrt{\Delta}} \tag{107}$$

$$p = .5 \tag{108}$$

These values will again satisfy the Martingale equality, the equality for the mean, and the variance of the S_i , in the same approximate sense.

7. Which Pricing Method to Choose?

In general, the choice of a pricing method depends on the following factors:

- The accuracy of pricing methods does, in general, differ. Some methods are numerically more stable than others. Some methods yield coarser approximations than others. Precision is an important factor.
- The speed of pricing methods also changes from one method to another. In general, everything else being the same, the faster results are preferred.
- Some methods are easier to implement. The ease of understanding a pricing method is an important factor in its selection by practitioners.

²³ This is, in fact, a standard assumption used throughout calculus. We notice that, as Δ goes to zero, the values of p will converge to $\frac{1}{2}$.

- The parsimony associated with the model is also important. In general, we want our pricing models to depend on as few parameters as possible. This way, the model has to be calibrated to a smaller number of parameters, which means that fewer things can go wrong. Also, a trader/broker can in general compensate for a parsimonious model by adjusting the quotes based on trading experience.

However, in the end, the method chosen depends on the circumstances, and is a matter of experience. What a book like this can do is to present a brief overview of the various approaches available to the financial engineer.

8. Conclusions

We obtained some important results in this chapter. First, we showed that the notion of state prices can be made practical in environments with liquid option prices at different strikes. From here we showed how to obtain risk-neutral and forward measures and the corresponding arbitrage-free dynamics.

Finally, as long as liquid option prices with different strikes exist, we showed how to replicate an asset using a static portfolio of options. This is true for the following reasons:

1. Given the option prices, we can get the prices of elementary insurance contracts.
2. But we know that every asset can be synthetically created as a portfolio of elementary insurance contracts.
3. This means that every asset can be created as a portfolio of liquid options.

Hence, option markets not only provide close relatives of elementary insurance contracts, but also show us how to obtain generalized static synthetics for all assets in principle. Of course, the practical application depends on the availability of liquid options.

Finally, we must emphasize that risk management and pricing are never as straightforward in real life, since given the day, the number, and the type of liquid option, contracts change.

Suggested Reading

*The treatment of the fundamental theorem of finance in this chapter has been heuristic and introductory, although all important aspects of the theorem have been covered. The reader can get more insight into the theorem by looking at **Duffie** (2001), which offers an excellent treatment of asset pricing. The article by **Brace et al.** (1997) is an important milestone in the use of Martingale theory, and places the right emphasis on pricing and the measure of change that fits this chapter. **Clewlow and Strickland** (1998) provide several examples.*

APPENDIX 11-1: Simple Economics of the Fundamental Theorem

This appendix provides a justification for the fundamental theorem from standard microeconomic theory. Consider the following setup. An investor faces a decision that involves two time periods; the time of decision, and T , the relevant future date. At T , there are only two possible states of the world $\omega^i, i = 1, 2$. The investor's subjective probabilities for these are p^1 and p^2 , respectively.

This investor's preferences are described by a *utility function* $U(X_t)$, where X_t is total (real) consumption at time t . Essentially, this investor is better off the higher his or her consumption:

$$0 < \frac{dU}{dX_t} \quad (109)$$

But, additional consumption would incrementally have less and less positive effect:

$$\frac{d^2U}{dX_t^2} < 0 \quad (110)$$

This investor would like to maximize the *expected utility* associated with his or her current and future consumption:

$$E_t^P[U(X_t) + \beta U(X_T)] = U(X_t) + \beta (p^1 U(X_T^1) + p^2 U(X_T^2)) \quad (111)$$

where β is a constant subjective discount factor, P is the subjective personal probability, and X_T^1 and X_T^2 are the consumption levels in states 1 and 2 during period T respectively. The maximization of this function is subject to the investor's budget constraint at time t and on the two states of the world at time T .

$$\begin{aligned} q_t X_t + S_t h_t &= I \\ q_T^1 X_T^1 &= I + h_t S_T^1 \\ q_T^2 X_T^2 &= I + h_t S_T^2 \end{aligned} \quad (112)$$

S_t is a risky asset that can be purchased at time t . It has possible values S_T^1 and S_T^2 at time T . Here the I is a known and constant income earned at times t and T . The q_t, q_T^1 , and q_T^2 are the corresponding prices of the consumption good. Note that, at time T , there are two prices, one for each state. Finally, h_t is the number of S_t purchased by the investor at time t .

According to this, we are dealing with an investor who receives a constant income that needs to be split between saving and consumption in a two-period setting. The investment can be made only by buying a desired amount of the S_t asset. The price of this asset is a random variable in the model.

The investor is risk averse and maximizes the expected utility function. There are several ways one can solve this maximization. Our intention is to show a simple example to *motivate* the fundamental theorem of finance. Hence, we are not concerned with the optimal consumption itself. Rather, we would like to obtain a relationship between "current" asset price S_t and the two possible values S_T^1 and S_T^2 at time T . The fundamental theorem of asset pricing is about these two sets of prices. Thus, we should be able to find out how the present framework can generate the state prices Q^i of the fundamental theorem.

Keeping these objectives in mind, we proceed by first substituting out the X_t, X_T^1, X_T^2 from the equations in (112), and then differentiating the resulting expression with respect to the only remaining choice variable h_t . The substitution gives

$$U(X_t) + \beta(p^1 U(X_T^1) + p^2 U(X_T^2)) = U\left(\frac{I - S_t h_t}{q_t}\right) + \beta\left(p^1 U\left(\frac{I + h_t S_T^1}{q_T^1}\right) + p^2 U\left(\frac{I + h_t S_T^2}{q_T^2}\right)\right) \quad (113)$$

Differentiating the right side with respect to h_t , equating to zero, and then rearranging,

$$U'\left(\frac{I - S_t h_t}{q_t}\right)\left(\frac{S_t}{q_t}\right) = \beta\left(p^1 U'\left(\frac{I + h_t S_T^1}{q_T^1}\right)\left(\frac{S_T^1}{q_T^1}\right) + p^2 U'\left(\frac{I + h_t S_T^2}{q_T^2}\right)\left(\frac{S_T^2}{q_T^2}\right)\right) \quad (114)$$

where the $U'(\cdot)$ is the derivative of the $U(x)$ with respect to “ x .”

Now comes the critical point. We can rearrange the first-order condition in equation (114) to obtain

$$S_t = \beta\left(p^1 \frac{U'\left(\frac{I + h_t S_T^1}{q_T^1}\right)}{U'\left(\frac{I - S_t h_t}{q_t}\right)} \frac{q_t}{q_T^1} S_T^1 + p^2 \frac{U'\left(\frac{I + h_t S_T^2}{q_T^2}\right)}{U'\left(\frac{I - S_t h_t}{q_t}\right)} \frac{q_t}{q_T^2} S_T^2\right) \quad (115)$$

Now relabel as follows:

$$Q^1 = \beta p^1 \frac{U'\left(\frac{I + h_t S_T^1}{q_T^1}\right)}{U'\left(\frac{I - S_t h_t}{q_t}\right)} \frac{q_t}{q_T^1} \quad (116)$$

and

$$Q^2 = \beta p^2 \frac{U'\left(\frac{I + h_t S_T^2}{q_T^2}\right)}{U'\left(\frac{I - S_t h_t}{q_t}\right)} \frac{q_t}{q_T^2} \quad (117)$$

It is clear that all elements of the right-side expressions are *positive* and, as a result, the Q^i , $i = 1, 2$ are positive. Substituting these Q^i back in equation (115), we get

$$S_t = S_T^1 Q^1 + S_T^2 Q^2 \quad (118)$$

In other words, there is a *linear* relationship between current asset price S_t and the future possible values S_T^1 and S_T^2 , and $\{Q^i\}$ is the determining factor.

An interesting implication of the derivation shown here is the following. Even when the utility function $U(\cdot)$ and the subjective probabilities p^i differ among investors, general equilibrium conditions would equate the marginal rates of substitution across these differing investors and hence the $\{Q^i\}$ would be the *same*. In other words, the $\{Q^i\}$ would be unique to all consumers even when these consumers disagree on the expected future behavior of the economy.

Exercises

1. The current time is $t = 1$ and our framework is the Libor model. We consider a situation with four states of the world ω_i at time $t = 3$.

Suppose L_i is the Libor process with a particular tenor and $B(1, 3)$, $B(1, 4)$, and $B(1, 5)$ are zero-coupon bond prices with indicated maturities. The *possible* payoffs of these instruments in the four future states of the world are as follows:

$$L = 6\%, 6\%, 4\%, 4\% \quad (119)$$

$$B(1, 3) = 1, 1, 1, 1 \quad (120)$$

$$B(1, 4) = 0.9, 0.92, 0.95, 0.96 \quad (121)$$

$$B(1, 5) = 0.8, 0.84, 0.85, 0.88 \quad (122)$$

The current prices are, respectively,

$$1, 0.91, 0.86, 0.77 \quad (123)$$

Here the 1 is a dollar invested in Libor. It is like a savings account. Finally, current Libor is 5%.

- Using Mathematica, determine a state price vector q_1, q_2, q_3, q_4 , that corresponds to $B(1, 3)$, $B(1, 4)$, $B(1, 5)$, L as a basis.
 - Does q_i satisfy the required condition of positivity? Is there an arbitrage opportunity?
 - Let F be the 1×2 FRA rate. Can you determine its arbitrage-free value?
 - Now let C be an ATM caplet (i.e., the strike is 5%) that expires at time $t = 2$, but settled at time $t = 3$ with notional amount 1. How much is it worth?
2. Suppose you are given the following data. The risk-free interest rate is 4%. The stock price follows:

$$dS_t = \mu S_t + \sigma S_t dW_t \quad (124)$$

The percentage annual volatility is 18% a year. The stock pays no dividends and the current stock price is 100.

Using these data, you are asked to calculate the current value of a European call option on the stock. The option has a strike price of 100 and a maturity of 200 days.

- Determine an appropriate time interval Δ , such that the binomial tree has five steps.
 - What would be the implied u and d ?
 - What is the implied “up” probability?
 - Determine the binomial tree for the stock price S_t .
 - Determine the tree for the call premium C_t .
3. Suppose the stock discussed in the previous exercise pays dividends. Assume all parameters are the same. Consider three forms of dividends paid by the firm:
- The stock pays a continuous, known stream of dividends at a rate of 4% per time.
 - The stock pays 5% of the value of the stock at the third node. No other dividends are paid.

- (c) The stock pays a \$5 dividend at the third node.

In each case, determine the tree for the ex-dividend stock price. For the first two cases, determine the premium of the call.

In what way(s) will the third type of dividend payment complicate the binomial tree?

4. We use binomial trees to value American-style options on the British pound. Assume that the British pound is currently worth \$1.40. Volatility is 20%. The current British risk-free rate is 6% and the U.S. risk-free rate is 3%. The put option has a strike price of \$1.50. It expires in 200 days.

- (a) The first issue to be settled is the role of U.S. and British interest rates. This option is being purchased in the United States, so the relevant risk-free rate is 3%. But British pounds can be used to earn the British risk-free rate. So this variable can be treated as a continuous rate of dividends. Or we can say that interest rate differentials are supposed to equal the expected appreciation of the currency.

Taking this into account, determine a Δ such that the binomial tree has five periods.

- (b) Determine the implied u and d and the relevant probabilities.
 (c) Determine the tree for the exchange rate.
 (d) Determine the tree for a European put with the same characteristics.
 (e) Determine the price of an American-style put with the previously stated properties.
5. Barrier options belong to one of four main categories. They can be up-and-out, down-and-out, up-and-in, or down-and-in. In each case, there is a specified “barrier,” and when the underlying asset price down or up-crosses this barrier, the option either expires automatically (the “out” case) or comes into life automatically (the “in” case).
- Consider a European-style up-and-out call written on a stock with a current price of 100 and a volatility of 30%. The stock pays no dividends and follows a geometric price process. The risk-free interest rate is 6% and the option matures in 200 days. The strike price is 110. Finally, the barrier is 120. If the before-maturity stock price exceeds 120, the option automatically expires.
- (a) Determine the relevant u and d and the corresponding probability.
 (b) Value a call with the same characteristics but without the barrier property.
 (c) Value the up-and-out call.
 (d) Which option is cheaper?